

6/pths

**SYSTEM AND METHOD FOR COMMUNICATING DATA BETWEEN  
NETWORKS OPERATING UNDER DIFFERENT PROTOCOLS**

The present invention relates to the field of communications in general and to communications systems supporting more than one communications protocol and methods relating thereto in particular.

5

The International Standard Organization (ISO) has defined a standard model for networking (or communications) protocols known as the Open System Interconnect (OSI) model. This model is based on seven protocol layers. Other models for layered protocols have been specified by the Internet Engineering Task Force (IETF), such as one based on Transmission Control  
10 Protocol / Internet Protocol (TCP/IP), see IETF publications RFC 793, RFC 791.

In a communications system supporting more than one communications protocol there is a need to transfer data between networks operating under the different protocols. A protocol conversion device, such as a Transport Service Bridge (TSB) or a Transport-Layer Gateway  
15 (TLG) can provide for such data transfer. This transfer is facilitated if the upper layers (e.g. session, presentation, and application entities) of the different protocols are practically identical, with only the lower layers differing. Certain current protocols exhibit such practically identical higher layers of protocol (e.g.. session, presentation and application layers). A protocol conversion device, such as a Transport Service Bridge (TSB) provides for  
20 communication between two networks operating under different protocols but where there is a commonality of protocol above a certain layer, e.g. layer-4.

Conversion is implemented in a so-called "buffer layer" which is used to convert the data formats between that expected by one protocol and that expected by the other. In a TCP/IP to OSI example, a TSB would make the TCP layer service look like an OSI TP4 layer service. This is accomplished by making the TCP layer emulate the TP4 layer. With such a TSB, the  
5 OSI application layer can extend over TCP/IP networks. Hence the same application layer can be used in both types of network. Only the underlying layers change between the two.

An example, by no means restricting the present invention, of a TSB is described in the IETF document RFC1006 "ISO Transport Service on top of the TCP" Version: 3, May 1987. RFC  
10 1006 defines how to produce OSI transport layer services on top of TCP and is incorporated herein by reference. Other types of TSB are described, for example, in other IETF RFCs.

A drawback of using a TSB is that it can produce a single point of failure in the communications path between any two terminals that communicate through it. The use of the  
15 term "terminals" here is to be understood to encompass "end systems" in an OSI network and "hosts" in a TCP/IP network. A single point of failure is a point of failure that cannot be avoided as there is no alternative path to allow data to bypass it. To avoid a TSB becoming a single point of failure, one or more additional TSBs can be connected in parallel with the first TSB (Figure 1).

20

Some protocols comprise both connectionless and connection-orientated layers. If we consider the example of a network operating under OSI protocol and one operating under TCP/IP, the layers 4 associated with each of the two different protocols will be connection-oriented

whereas the layer below (i.e. layer-3 - the Network layer where routing is carried out) will be connectionless. The connectionless nature of the Network layers means that, if more than one TSB are provided in parallel, it will be necessary to ensure that all messages communicated between a particular pair of terminals pass through the same TSB. If this is not done, the connection oriented nature of the layer-4 causes the connections to be unstable. An appropriate control protocol is therefore required to nominate one of the TSBs as the "active" TSB used to provide the communications path for those two terminals.

This instability can be illustrated using the analogy of a conventional telephone call. A connection-oriented layer-4 connection is analogous to a phone call in that there is a 'dialling' (set-up) phase, a 'call' phase, and a 'clear-down' phase. For the purposes of this analogy, the dialling phase of the call includes the transfer of further information beyond simple addressing information (i.e. agreements on how to govern the conversation – how long each party can speak without an acknowledgement, etc.). During the call phase, the conversation can be regarded as the exchange of a number of packets of data. However, unlike the words in an ordinary phone call, each packet has sufficient addressing information to get it to its destination, and contains further information such as a sequence number and a specific reference number for that call.

If party A and party B are situated in different administrations, they will have to go through a "bridge" function that is set up when the call is established. The call in each network has to be set up separately and a connection established between the two networks in the bridge. The directory number (address) used by party A in network A to indicate party B cannot be used in network B: it has to be translated in the bridge to a new format that is understood by network

B. If we imagine a situation where two paths exist between the two parties (party A and party B) that wish to converse, each path may pass via a different one of the bridges.

If party A calls party B and the call is established via bridge 1, two connections are used: one  
5 between party A and bridge 1 and one between bridge 1 and party B. Bridge 1 and bridge 2  
both have routes to party B and although party A has established a connection with bridge 1, as  
mentioned earlier, each packet has complete addressing information and there is the possibility  
that some packets could be routed to bridge 2. As no connection exists with bridge 2 these  
packets will get discarded, interrupting the conversation. As a result, the connection with  
10 bridge 1 may also clear down (timeout).

The present invention provides a communications system comprising a first network for the  
communication of data according to a first protocol and a second network for the  
communication of data according to a second protocol; in which each network comprises at  
15 least one node, in which the system also comprises a plurality of communication interfaces for  
providing communication between a first node of the first network and a second node of the  
second network; in which each interface comprises means for sending values to the first node  
for indicating the availability of communication between that interface and the second node; in  
which the system comprises selection means for selecting one of the interfaces for  
20 communicating data between the first node and the second node based on the values sent by the  
interfaces to the first node in which each interface comprises means for detecting the selection  
of one of the plurality of interfaces and means for modifying, on selection of the selected  
interface, the value sent to the first node in which the selection means comprises means for  
preferentially selecting an interface associated with the modified value.

The present invention also provides a method for the communication of data between a first node in a first network and a second node in a second network, in which data is communicated in the first network according to a first protocol and data is communicated in the second network according to a second protocol; in which the method includes the step of providing a plurality of communication interfaces for providing communication between the first node and the second node; in which the method also includes the step of sending values from each interface to the first node indicating the availability of communication between that interface and the second node; in which the method also includes the step of selecting one of the interfaces for communicating data between the first and second nodes, in which the selection is based on the values sent by the interfaces to the first node including the steps of detecting the selection of the selected interface and modifying, the value sent by the selected interface to the first node including the step of preferentially selecting the interface sending the modified value.

15

Embodiments of the invention will now be described, by way of example only, with reference to the drawings, in which:

Figure 1 shows a conventional communications system using TSBs;

20

Figures 2 and 3 show the protocol layers of the system of Figure 1;

Figure 4 shows a representation of information flow in a TSB according to an embodiment of the present invention;

Figures 5 to 10 show operation of two TSBs according to further embodiments of the present invention.

Figure 1 shows a TCP/IP network connected to an OSI network via a number of TSBs connected in parallel (only two TSBs are shown for simplicity). Each of the two networks will typically have a large number of attached terminals, one being shown per network for the purposes of illustration. For the purposes of the present description, the terminal connected to the TCP/IP network is designated "source" and the terminal connected to the OSI network is designated "destination".

Figure 1 shows a conventional arrangement in which two parallel TSBs are used to provide communication between different networks: in this example between a first network operating under the TCP/IP protocol and a second network operating under the OSI protocol. In order to provide this communication, the TSBs act to interface between the lower layers of the two protocols and to transfer the upper layers of a first protocol (e.g. OSI) from the "non-native" lower layers of a second protocol (e.g. TCP/IP) to the "native" lower layers of its own, first protocol. If we look at the end-to-end path, a common set of upper, OSI layers is carried across a region of "non-native" lower layers. The corresponding protocol layers are shown in Figure 2.

Figure 2 illustrates the structure of the protocol layers at the source and destination terminals of Figure 1 and in addition the two sets (or "stacks") of protocol layers of one of the TSBs linking the two networks. Figure 2 illustrates the source terminal, operating in the TCP/IP network that supports a protocol including a number of upper layers, not shown in detail, of OSI protocol. A layer according to RFC 1006 provides accommodation between the upper, OSI,

BEST AVAILABLE COPY

layers and the lower TCP layer which itself is above an IP layer. Further, layers below the IP layer (such as Ethernet or PPP for IP, IEEE 8802.2, 8802.3 for OSI over Ethernet and Q.921 for OSI over SDH DCC channels) are not shown in detail. The source terminal is connected to the TSB via a TCP/IP Data Communications Network (DCN). The TSB supports a TCP/IP-compatible protocol stack comprising similar layers to that of the source terminal. In addition the TSB supports an OSI-compatible protocol stack, the lower layers of which comprise a Transport Protocol Class 4 (TP4) layer above a Connectionless Network Layer Protocol (CLNP) layer. These layers are defined as follows: TP4 in ISO 8072, ISO 8073, ITU-T X.214, ITU-T X.224 and CLNP in ISO 8473. Upper layers, i.e. above TCP or TP4 are not terminated by the TSB. Instead, a single TSB Interworking Function joins the two different protocol stacks together. The TSB communicates with the destination terminal via an OSI DCN. The destination terminal supports OSI protocol layers with a number of upper OSI protocol layers, not shown in detail, above a TP4 layer which itself is above a CLNP layer.

Figure 3 shows the same network as Figure 2 but illustrating communication between a source terminal in an OSI DCN supporting OSI protocol upper layers above OSI lower protocol layers and a destination terminal in a TCP/IP DCN supporting OSI protocol upper layers above TCP/IP lower protocol layers.

Advantageously, the present invention also applies to communication of TCP/IP protocol upper layers, such as hypertext transfer protocol (HTTP), Telnet, etc across TCP/IP and OSI DCNs. In this case, the RFC 1006 layer is replaced by an equivalent layer able to accommodate TCP/IP upper layer over OSI lower layers.

The present invention advantageously provides for the active TSB to be chosen by the source terminal from a group of TSBs, when initially setting up the communications path, and advantageously copes with failures in either network or in a TSB of a group of TSBs

5 In operation, TSBs advertise their availability by sending messages to connected terminals. According to a preferred embodiment, the present invention uses information from both protocol layers 3 (network/routing) and 4 (transport) to control the advertising of availability by TSBs. The basic control structure according to this preferred embodiment is shown in Figure 4. Routing information is exchanged (see arrows 1) between the TSB and the terminals  
10 in each network using level 3 of the appropriate protocol, i.e. IP routing for the TCP/IP DCN and OSI routing for the OSI DCN. Figure 4 shows the protocol layers of a TSB with the TSB interworking function in more detail. As shown in Figure 4, the TSB interworking function comprises functional blocks for connection-control, address-conversion and metric-control. The connection-control block monitors the connection state of the layer-4 connection on each  
15 side of the TSB and co-ordinates the set-up and teardown of the link. The address-conversion block controls the conversion between the addressing used by the lower layers of the protocol on one side of the bridge (IP address/port number) and that used by the equivalent layers on the other side (network service access point/ transport service access point: NSAP/TSAP). The mapping between addresses may be set up either by a protocol, manually, or algorithmically  
20 (i.e. by a conversion function which processes the address of one domain (or DCN) according to a set of rules and consequently forms an address which is acceptable to the other domain (or DCN) ). The metric control block controls which metric values are associated with each route advertised by the routing protocols. Here a metric is a value used in assessing the desirability of one route over another, it can be thought of as a cost..



The arrows 2 to 5 in Figure 4 show the flow of information within the TSB. Status control arrows 2 indicate communication, from layer 3 of the appropriate protocol to the address conversion block, of the availability of routes in the relevant domain (i.e. OSI DCN or TCP/IP DCN). In the OSI domain, routing is controlled by an IS-IS routing function within layer-3. Here, IS refers to "intermediate system". This availability information indicates whether to advertise each address. The address conversion block in turn provides to the metric control block the OSI address information 3 for use in the advertising of routes in the IP domain by the metric control block (or IP address information for the OSI domain if working the other way round). Metric control arrow 4 indicates communication of information from the connection control block to the metric control block to allow link failures in either domain to also influence the advertising of routes by the metric control block by means of the metrics associated with each route. The metric control block then communicates (see arrows 5) to the respective routing protocol the metric values to use in the appropriate DCN.

Hence routes are advertised by combining a metric value with an address value. Layer 3 of the TCP/IP protocol (the IP layer) controls the advertising of the IP routes and layer 4 (the TCP layer) controls the metric value used (e.g. a default value or a lower value).

The terminals choose the "active" TSB using the routing process. For example, in the TCP/IP domain, the terminal will use the TSB that advertises a route to the required destination with the lowest metric value. Two transport connections (layers-4 of the two protocols) are established by the terminals and the active TSB (one connection from the source to the active TSB and one from the active TSB to the destination). Once these connections are established,

the routing is "locked" to the active TSB by modification of the route metrics advertised by the active TSB. If the "active" TSB detects a failure affecting the chosen route, the TSB stops advertising the route until availability can again be offered. Availability is offered again by re-advertising the route availability with the original, default metric value.

5

Figure 5 shows the two TSBs of Figure 1 advertising to the source terminal in the TCP/IP network their availability and the availability of a route from that TSB to the destination in the OSI network prior to the setting up of a connection from the source to the destination. Here both TSBs have identified a route to the destination and, as neither has been selected by the source, each is advertising availability associated with the same (i.e. the default) metric value. The TSBs map each IP address in the TCP/IP network relating to a destination in the OSI network onto a different NSAP address in the OSI network. Each of the TSBs advertises its availability for a particular IP address relating to a destination in the OSI network if a route is available from that TSB to the corresponding destination in the OSI network and if an address conversion (OSI to TCP/IP) is available. OSI to TCP/IP address conversion may be by table look-up or algorithmic

10  
15

20

Each TSB advertising the availability of a route to a particular destination in the OSI network has the ability to terminate the IP connection as if that TSB was the destination IP address. This means that multiple instances of the same TCP/IP terminal will effectively exist in the TCP/IP network. At first this seems to break the rules of IP network addressing as each TCP/IP terminal should be uniquely addressable. However, a routing function is included in the TSBs so that the destination IP address(es) do not map directly to physical ports but belong to virtual ports supported by the routing function. This means that the IP addresses are not directly

accessible at the physical interfaces of the TSB (they are only accessible through the TSB routing function) and hence multiple TSBs can, as a result, be directly connected to the same physical network (i.e. IP sub-net).

- 5 As indicated above, significant problems arise if the route between a source and destination changes. This is because TCP is connection-oriented and relies on sequential delivery of packets from the protocol layer-3 (IP layer). A change in routing would break the sequential delivery of packets from the IP layer and cause the original TCP connection to be terminated. This would in turn cause the OSI connection to be terminated. Furthermore, if a new route is
- 10 established via a different TSB, a new TCP connection will not be established immediately as the source will first have to recognise that the connection has been interrupted and clear itself back to a state from which it can reattempt a connection.

When the TCP connection is established on the new TSB there may still be difficulties in

15 establishing a new OSI connection to the destination as only one connection is allowed into the destination at any time and the previous OSI connection may not have timed-out. In order to avoid unwanted changes of route and the consequent problems outlined above, the invention provides for the connection-state of the OSI connection to influence the routing information held in the TSB such that if a TSB has a connection established for a certain IP address, it

20 changes the metric associated with the associated route to a lower value (see State 2 in Figure 6). Figure 6 shows the communication system of Figure 5 following establishment of a TCP/IP connection from the source to TSB1 and the establishment of an OSI connection from TSB1 to the destination. No connection between the destination and source is established via TSB2 which is, as a result, still advertising availability using the default metric. TSB1 has a

connection established through it, and is advertising availability with a lower value metric. This effectively “locks” the routing to the chosen TSB (i.e. TSB1).

Figures 7 to 10 illustrate how the system of Figure 5 copes with a failure condition.

5 Advantageously, the advertising of multiple available routes via parallel TSBs allows fast recovery from failure conditions, as follows. Failure of either TCP/IP or OSI connections used by an active TSB will cause the active TSB to cease advertising availability. Advertising will either cease indirectly (i.e. if the connection fails in the TCP/IP network) or directly (i.e. if the connection fails in the OSI network): see States 3 – 6 in Figures 7 to 10.

10

Figures 7-10 show the communication system of Figure 5 in the presence of a fault on the OSI connection from the destination terminal to the active TSB: TSB1. TSB1 reacts to detection of this fault by ceasing to advertise availability to the source. TSB2 is unaffected by the fault and continues to advertise availability in association with the default metric. Figure 8 shows the source reacting to the lack of advertising of availability by TSB1 by choosing TSB2 to connect to the destination. TSB2, now selected, changes the metric from the default value to a low value. Figure 9 shows the situation some time later when TSB1 has re-established availability of a path to the destination and has recommenced advertising availability to the source using the default metric. Normally there would be no reason for the source to change the TSB it was using to connect with the destination as TSB2 is still advertising availability using the lower metric and would therefore continue to be preferred by the source. However, Figure 9 also shows a fault in the TCP connection from TSB2 to the source. In Figure 10 we see the result of the source terminal detecting this fault and changing to TSB1 for connection to the destination. On being selected by the source terminal, TSB1 continues to advertise availability but now in

15

20

association with a lower metric. TSB2 ceases to advertise availability on sensing loss of communication with the source terminal. Once the failure condition is removed TSB2 may re-advertise availability.

- 5 Advantageously, the present invention provides simple operation with minimal interaction between TSBs (i.e. avoiding a complicated "hot standby" protocol), whilst exploiting the facilities of existing protocols and supporting single or multiple backup TSBs.

The invention is not limited to the embodiments set out above which are used to illustrate the  
10 principle thereof. In particular, the routing may equally be "locked" to the active TSB by modification of the route metrics advertised by the non-active TSB(s). The invention has application to instances of 3 or more TSBs arranged effectively in parallel.

The present invention is also applicable to alternative forms of interface for providing  
15 communications between networks operating under different protocols, including but not limited to Transport-Layer Gateways.

The invention has been mainly described, by way of example, with reference to a source terminal in a TCP/IP network and a destination terminal in an OSI network. As indicated  
20 above, the invention also applies to systems in which the source entity is an OSI terminal and the destination is a TCP/IP terminal and to networks operating according to other protocols..